

Exhibit V



PATENTS
112056-0126
P01-1108

Honorable Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

RECEIVED
AUG 27 2004
Technology Center 2100

AMENDMENT TRANSMITTAL LETTER

Sir:

Transmitted herewith for filing in the Patent Application of

Inventor(s): Steven R. Kleiman et al.

For: A METHOD FOR WRITING CONTIGUOUS ARRAYS
OF STRIPES IN A RAID STORAGE SYSTEM
USING ASSOCIATED BLOCK WRITES

Serial No.: 10/105,034

Filed: March 21, 2002

are the following papers:

- ☒ A response to the Office action dated *May 20, 2004*
- ☐ A petition for extension of time and check for \$
- ☒ A check for \$618 to cover the fee for additional claims resulting from the response. The fee is calculated as follows:

PATENTS
112056-0126
P01-1108

Total Claims Remaining after
Amendment:.....64
Less Highest Number of Total Claims
Previously Paid for:.....44

Difference.....20 × \$ 18 = \$ 360

Independent Claims Remaining
after Amendment.....13
Less Highest Number of Independent
Claims Previously Paid for.....10

Difference.....3 × \$ 86 = \$ 258

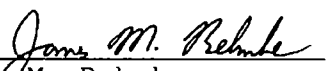
Multiple-Dependent-Claim Fee Required
after Amendment and Not Previously Paid..... \$ 0

Total of Above Calculations..... \$ 618
Less Reduction for Small Entity \$ 0

Total Due \$ 618

The Commissioner is hereby authorized to charge any
other fees under 37 C.F.R. §§1.16 and 1.17 that may be required,
or credit any overpayment, to our Deposit Account No. 03-1237.

Respectfully submitted,


James M. Behmke
Reg. No. 51,448
CESARI AND MCKENNA, LLP
88 Black Falcon Avenue
Boston, MA 02210
(617) 951-2500

PATENTS
112056-0126
P01-1108

IN THE TITLE:

Please replace the previous title with the following new title:

- A METHOD FOR WRITING CONTIGUOUS ARRAYS OF STRIPES IN A
RAID STORAGE SYSTEM USING ASSOCIATED BLOCK WRITES -

PATENTS
112056-0126
P01-1108

IN THE CLAIMS:

Please re-write the claims to read as follows:

- 1 1. - 16. (Cancelled)
- 2 17. (Original) A method for controlling storage of data in a plurality of storage devices each
3 comprising storage blocks, the method comprising:
- 4 buffering a plurality of write requests associated with data blocks for a single write
5 transaction;
- 6 defining a group of storage blocks, the group comprising a plurality of storage blocks
7 in each of the plurality of storage devices; and
- 8 associating each data block with a respective one of the storage blocks, for transmit-
9 ting the association to a storage device manager for processing of the single write transaction.
- 1 18. (Original) The method of claim 17, further comprising receiving by a file system the
2 plurality of write requests, and transmitting from the file system to the storage device man-
3 ager the association for processing of the single write transaction, wherein the storage device
4 manager comprises a RAID layer.
- 1 19. (Original) The method of claim 17, wherein associating comprises associating each data
2 block of at least one of the write requests with storage blocks of only one of the plurality of
3 storage devices.

PATENTS
112056-0126
P01-1108

1 20. (Original) The method of claim 17, wherein defining the group of storage blocks com-
2 prises receiving by the file system from the storage device manager configuration informa-
3 tion of the plurality of storage devices.

1 21. (Original) The method of claim 20, wherein the configuration information comprises
2 disk topology information.

1 22. (Original) The method of claim 17, wherein the group of storage blocks comprises more
2 than one stripe in a RAID group.

1 23. (Original) The method of claim 22, wherein the group of storage blocks comprises more
2 than one contiguous stripe.

1 24. (Original) The method of claim 23, further comprising transmitting the association and
2 the data to the plurality of storage devices to store each data block at the respective one of the
3 storage blocks of the more than one contiguous stripe.

1 25. (Original) The method of claim 22, wherein the plurality of storage devices comprises
2 more than one RAID group, and defining the group of storage blocks comprises defining the
3 group of storage blocks from storage blocks in a first RAID group.

1 26. (Original) The method of claim 25, further comprising defining a second group of stor-
2 age blocks in a second RAID group for association with data associated with a second plu-
3 rality of write requests for processing of a second write transaction.

PATENTS
112056-0126
P01-1108

- 1 27. (Original) The method of claim 17, further comprising queuing a plurality of associa-
2 tions by the storage device manager for processing of a plurality of write transactions.
- 1 28. (Original) The method of claim 27, further comprising processing the plurality of write
2 transactions by the storage manager in an order derived from a plurality of priorities associ-
3 ated with the plurality of write transactions.
- 1 29. (Original) The method of claim 17, wherein the group of storage blocks comprises a
2 same quantity of storage blocks in each one of the plurality of storage devices.
- 1 30. (Original) The method of claim 29, wherein each of the plurality of storage blocks com-
2 prises contiguous storage blocks.
- 1 31. (Original) The method of claim 17, wherein the group of storage blocks comprises allo-
2 cated and unallocated storage blocks.
- 1 32. (Original) The method of claim 31, wherein defining the group further comprises se-
2 lecting an unallocated storage block having a lowest block identifier of any unallocated stor-
3 age block as a storage block having a lowest storage block identifier of any storage block in
4 the group.
- 1 33. (Original) The method of claim 17, wherein the write requests comprise data to be
2 written.
- 1 34. (Original) A method for storing data blocks, the method comprising:

PATENTS
112056-0126
P01-1108

2 providing a RAID layer in communication with a plurality of storage devices that
3 each comprise a plurality of storage blocks;

4 receiving by the RAID layer a write transaction request that includes an association of
5 each data block with a respective one of a group of storage blocks, the group comprising a
6 plurality of storage blocks in each of the plurality of storage devices; and

7 storing the data blocks by the RAID layer in the group of storage blocks according to
8 the association.

1 35. (Original) The method of claim 34, wherein the RAID layer comprises a second plural-
2 ity of storage devices that each comprise a plurality of storage blocks that mirror storage
3 blocks of the plurality of storage devices, and further comprising storing the data in a group
4 of storage blocks of the second plurality of storage devices.

1 36. (Original) The method of claim 35, further comprising performing a parity determina-
2 tion for one of the groups of storage blocks prior to storing, and sharing a result of the parity
3 determination with the other of the groups of storage blocks.

1 37. (Original) The method of claim 36, further comprising selecting one of the groups of
2 storage blocks for performing the parity determination.

1 38. (Original) The method of claim 34, further comprising monitoring by the storage device
2 manager of configuration information of the plurality of storage devices.

1 39. - 41. (Cancelled)

PATENTS
112056-0126
P01-1108

1 42. (Original) A device for controlling storage of data in a plurality of storage devices each
2 comprising storage blocks, the device comprising:

3 a buffer collecting write requests, each request associated with data blocks;

4 a processor defining a group of storage blocks, the group comprising a plurality of
5 storage blocks in each of the plurality of storage devices, the processor associating each data
6 block with a respective one of the storage blocks for a single write transaction; and

7 a transmitter transmitting the association to a storage device manager for processing
8 of the single write transaction.

1 43. (Original) An apparatus for storing data blocks, the apparatus comprising:

2 a plurality of storage devices that each comprise a plurality of storage blocks; and

3 a storage device manager in communication with the plurality of storage devices, and
4 configured to receive a write transaction from a file system, the write transaction comprising
5 an association of each data block to a respective storage block of a group of storage blocks,
6 the group comprising a plurality of storage blocks in each of the plurality of storage devices.

1 44. (Original) The apparatus of claim 43, wherein the storage device manager causes a par-
2 ity determination for the group of storage blocks prior to storing the data blocks in the group
3 of storage blocks.

PATENTS
112056-0126
P01-1108

Please insert the following new claims 45 *et seq.*:

- 1 45. (New) A device for controlling storage of data in a plurality of storage devices each
2 comprising storage blocks, the device comprising:
- 3 means for buffering a plurality of write requests associated with data blocks for a sin-
4 gle write transaction;
- 5 means for defining a group of storage blocks, the group comprising a plurality of
6 storage blocks in each of the plurality of storage devices; and
- 7 means for associating each data block with a respective one of the storage blocks, for
8 transmitting the association to a storage device manager for processing of the single write
9 transaction.
- 1 46. (New) A computer readable media, comprising: the computer readable media containing
2 instructions for execution in a processor for the practice of the method of,
- 3 buffering a plurality of write requests associated with data blocks for a single write
4 transaction;
- 5 defining a group of storage blocks, the group comprising a plurality of storage blocks
6 in each of a plurality of storage devices; and
- 7 associating each data block with a respective one of the storage blocks, for transmit-
8 ting the association to a storage device manager for processing of the single write transaction.
- 1 47. (New) Electromagnetic signals propagating on a computer network, comprising: the
2 electromagnetic signals carrying instructions for execution in a processor for the practice of
3 the method of,

PATENTS
112056-0126
P01-1108

4 buffering a plurality of write requests associated with data blocks for a single write
5 transaction;

6 defining a group of storage blocks, the group comprising a plurality of storage blocks
7 in each of a plurality of storage devices; and

8 associating each data block with a respective one of the storage blocks, for transmit-
9 ting the association to a storage device manager for processing of the single write transaction.

1 48. (New) An apparatus for storing data blocks, the apparatus comprising:

2 means for providing a RAID layer in communication with a plurality of storage de-
3 vices that each comprise a plurality of storage blocks;

4 means for receiving by the RAID layer a write transaction request that includes an as-
5 sociation of each data block with a respective one of a group of storage blocks, the group
6 comprising a plurality of storage blocks in each of the plurality of storage devices; and

7 means for storing the data blocks by the RAID layer in the group of storage blocks
8 according to the association.

1 49. (New) A computer readable media for use with storing data blocks, comprising: the
2 computer readable media containing instructions for execution in a processor for the practice
3 of the method of,

4 providing a RAID layer in communication with a plurality of storage devices that
5 each comprise a plurality of storage blocks;

6 receiving by the RAID layer a write transaction request that includes an association of
7 each data block with a respective one of a group of storage blocks, the group comprising a
8 plurality of storage blocks in each of the plurality of storage devices; and

PATENTS
112056-0126
P01-1108

9 storing the data blocks by the RAID layer in the group of storage blocks according to
10 the association.

1 50. (New) Electromagnetic signals propagating on a computer network for use with storing
2 data blocks, comprising: the electromagnetic signals carrying instructions for execution in a
3 processor for the practice of the method of,

4 providing a RAID layer in communication with a plurality of storage devices that
5 each comprise a plurality of storage blocks;

6 receiving by the RAID layer a write transaction request that includes an association of
7 each data block with a respective one of a group of storage blocks, the group comprising a
8 plurality of storage blocks in each of the plurality of storage devices; and

9 storing the data blocks by the RAID layer in the group of storage blocks according to
10 the association.

1 51. (New) A method for controlling storage of data, comprising:

2 receiving one or more write requests associated with data blocks;

3 receiving topological information associated with storage blocks of a storage system;

4 associating the data blocks with one or more storage blocks; and

5 writing the data blocks, in response to the association, to the one or more storage de-
6 vices in a single write request.

1 52. (New) The method of claim 51, further comprising: transmitting the association to a
2 storage device manager.

PATENTS
112056-0126
P01-1108

1 53. (New) The method of claim 51, further comprising: associating each data block with a
2 single storage block.

1 54. (New) The method of claim 51, further comprising: storing the data blocks in the asso-
2 ciation.

1 55. (New) The method of claim 51, further comprising: storing the data blocks in a memory
2 of the storage system.

1 56. (New) The method of claim 51, further comprising: creating an array as the association.

1 57. (New) The method of claim 51, further comprising: buffering a plurality of write re-
2 quests into the single write request.

1 58. (New) The method of claim 57, further comprising: buffering, in a buffer, the plurality
2 of write requests into the single write request until a predetermined criteria is met.

1 59. (New) The method of claim 58, further comprising: meeting the criteria when a buffer is
2 full.

1 60. (New) The method of claim 58, further comprising: meeting the criteria when the single
2 write request is a predetermined logical length.

PATENTS
112056-0126
P01-1108

1 61. (New) The method of claim 51, further comprising: receiving topological information of
2 a plurality of storage devices of the storage system.

1 62. (New) The method of claim 61, further comprising: using a RAID system as the plural-
2 ity of storage devices.

1 63. (New) The method of claim 51, further comprising: using a plurality of disks for the
2 storage system.

1 64. (New) The method of claim 51, further comprising: transmitting the association to the
2 storage device manager. ~

1 65. (New) The method of claim 51, further comprising: organizing the association as a
2 combination of columns and rows.

1 66. (New) The method of claim 65, further comprising: associating each column with a
2 storage device.

1 67. (New) The method of claim 66, further comprising: associating each row with a stripe
2 of the storage device.

1 68. (New) The method of claim 67, further comprising: arranging the association as a plu-
2 rality of stripes on a plurality of storage devices.

PATENTS
112056-0126
P01-1108

1 69. (New) A storage system, comprising:
2 a file system, the file system to receive one of more write requests associated with
3 data blocks;
4 a storage device manager, the storage device manager to generate topological infor-
5 mation of storage blocks of one or more storage devices, and to send the topological infor-
6 mation to the file system; and
7 an association generated in the file system, the association to associate the data blocks
8 with one or more storage blocks of the one or more storage devices, the association to be sent
9 to the storage device manager, the storage device manager to write the data blocks, in re-
10 sponse to the association, to the one or more storage blocks as a single write request.

1 70. (New) The storage system of claim 69, further comprising: an association that associates
2 each data block with a single storage block.

1 71. (New) The storage system of claim 69, further comprising: a memory to buffer the data
2 blocks for the write request to the one or more storage devices.

1 72. (New) The storage system of claim 69, further comprising: a memory to store the asso-
2 ciation containing the data blocks.

1 73. (New) The storage system of claim 69, further comprising: one or more storage devices
2 having storage blocks.

1 74. (New) The storage system of claim 69, further comprising: an array as the association.

PATENTS
112056-0126
P01-1108

1 75. (New) The storage system of claim 69, further comprising: a buffer in the file system to
2 receive the one or more write requests.

1 76. (New) The storage system of claim 69, further comprising: an arrangement of columns
2 and rows in the association.

1 77. (New) The storage system of claim 76, further comprising: each column representing a
2 storage device of the one or more storage devices.

1 78. (New) The storage system of claim 76, further comprising: each row representing a
2 stripe of the one or more storage devices.

1 79. (New) The storage system of claim 69, further comprising: a plurality of disks as the
2 one or more storage devices.

1 80. (New) The storage system of claim 69, further comprising: a RAID system as the plu-
2 rality of storage devices.

1 81. (New) A storage system, comprising:

2 means for receiving one or more write requests associated with data blocks;

3 means for receiving topological information associated with storage blocks of a stor-
4 age system;

5 means for associating the data blocks with one or more storage blocks; and

PATENTS
112056-0126
P01-1108

6 means for writing the data blocks, in response to the association, to the one or more
7 storage devices in a single write request.

1 82. (New) A computer readable media, comprising: the computer readable media containing
2 instructions for execution in a processor for the practice of the method of,
3 receiving one or more write requests associated with data blocks;
4 receiving topological information associated with storage blocks of a storage system;
5 associating the data blocks with one or more storage blocks; and
6 writing the data blocks, in response to the association, to the one or more storage de-
7 vices in a single write request.

1 83. (New) Electromagnetic signals propagating on a computer network, comprising: the
2 electromagnetic signals carrying instructions for execution in a processor for the practice of
3 the method of,
4 receiving one or more write requests associated with data blocks;
5 receiving topological information associated with storage blocks of a storage system;
6 associating the data blocks with one or more storage blocks; and
7 writing the data blocks, in response to the association, to the one or more storage de-
8 vices in a single write request.

PATENTS
112056-0126
P01-1108

REMARKS

This Amendment is filed in response to the Office Action mailed on May 20, 2004.
All objections and rejections are respectfully traversed.

Claims 17-38, and 42-83 are in the case.

Claims 45-83 have been added to better claim the invention.

Applicant's undersigned attorney would like to take this opportunity to thank Examiner Nguyen for the faxed copy of the above-mentioned Office Action.

At paragraph 3 of the Office Action, the title of the invention was objected to as being non-descriptive. A new title has been provided, and the title is believed to be in allowable condition.

At paragraph 4 of the Office Action, Examiner points out that the reference "DAVID HITZ et al" on the IDS filed 5/22/03 was not considered because it lacks a date. Although no publication date is immediately available, Applicant directs Examiner to the Manual of Patent Examining Procedure (MPEP) section 2129, which states, "When applicant states that something is prior art, it is taken as being available prior art against the claims." Accordingly, Applicant respectfully requests that the "DAVID HITZ et al" reference be considered prior art for this current application.

At paragraph 6 of the Office Action, claims 17-38, and 42-44 were rejected under 35 U.S.C. §102 (b) as being anticipated by DeKoning, U.S. Patent No. 6,148,368, issued on November 14, 2000.

PATENTS
112056-0126
P01-1108

The present invention, as set forth in representative claim 17, comprises in part:

A method for controlling storage of data in a plurality of storage devices each comprising storage blocks, the method comprising:

buffering a plurality of write requests associated with data blocks for a single write transaction;

defining a group of storage blocks, the group comprising a plurality of storage blocks in each of the plurality of storage devices; and

associating each data block with a respective one of the storage blocks, for transmitting the association to a storage device manager for processing of the single write transaction.

DeKoning discloses a method for accelerating disk array write operations using segmented cache memory and data logging. In DeKoning, write commands are first buffered to the cache segments, and later accumulated into a cache extension disk region at a later time, such as when the disks are idle. The accumulated write requests are the write commands combined into “larger write requests (e.g. RAID stripe writes)” (Col. 8, Lines 4-7, emphasis added) to spare the storage system from the write penalty of multiple smaller writes.

Applicant respectfully urges that DeKoning does not show Applicant’s claimed novel, “***associating each data block with a respective one of the storage blocks, for transmitting the association to a storage device manager for processing of the single write transaction.***”

Applicant’s claimed invention is directed toward the mapping of data blocks with the storage blocks to which they will be written. Applicant claims buffering a plurality of write requests and combining them into a single write transaction. In addition to this, however, Applicant goes one step further by associating each data block of the single write request with a storage block of the storage system before transmitting the buffered write request to a

PATENTS
112056-0126
P01-1108

storage device manager. In this way, the storage device manager is not required to map each data block to a storage block of the storage system, as it would have been required to do with a standard RAID stripe write command. Also, Applicant's claimed invention is not limited to sending single stripe write transactions, but can send write requests as associated blocks spanning multiple stripes of the storage devices. DeKoning does not address associating data blocks with storage blocks, but instead merely discusses the use of buffering smaller write requests into a larger write request, such as a RAID stripe write.

In summary, DeKoning combines multiple write requests into a single ordinary stripe write request. In sharp contrast, Applicant claims a further step, associating the multiple write requests with the actual storage blocks on the disks to which they are to be written.

Applicant respectfully urges that the DeKoning patent is legally precluded from anticipating the claimed invention under 35 U.S.C. §102 because of the absence from the DeKoning patent of Applicant's "*associating each data block with a respective one of the storage blocks, for transmitting the association to a storage device manager for processing of the single write transaction.*"

All independent claims are believed to be in condition for allowance.

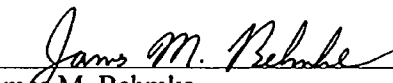
All dependent claims are believed to be dependent from allowable independent claims, and therefore in condition for allowance.

Favorable action is respectfully solicited.

PATENTS
112056-0126
P01-1108

Please charge any additional fee occasioned by this paper to our Deposit Account
No. 03-1237.

Respectfully submitted,


James M. Behmke
Reg. No. 51,448
CESARI AND MCKENNA, LLP
88 Black Falcon Avenue
Boston, MA 02210-2414
(617) 951-2500